Living Inside Arithmetic: Science, Technology, and Thought in a World Limited to Peano Axioms

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

December 9, 2025

Imagine a civilization whose entire formal understanding of reality is constrained to one modest theory: first-order Peano arithmetic. Its basic objects are the natural numbers; its primitive notions are zero, successor, addition, multiplication, and order. No sets, no uncountable hierarchies, no large cardinals—unless they can be coded and justified arithmetically. At first glance this sounds brutally restrictive, as if higher mathematics, theoretical physics, and sophisticated computation would all be suffocated at the foundation. Yet the shock is in the opposite direction. Peano arithmetic is remarkably expressive. It can formalize Turing machines, capture ordinary algorithms, support the mathematics of finite and countable structures, and provide enough real analysis for most of contemporary physics and engineering as actually practiced. Bridges still stand, rockets still launch, and quantum devices still operate. The real constraints appear not in day-to-day technical work, but in the deep structure of proof: certain true combinatorial principles remain unreachable, some program properties can never be certified, and large swathes of higher set theory become invisible. This paper explores how much a PA-limited world can accomplish, where the invisible ceilings lie, and what it would mean—epistemically and philosophically—for an entire scientific culture to live inside arithmetic.

Keywords: Peano arithmetic, formal systems, incompleteness, computability, reverse mathematics, physics foundations, verification limits, arithmetic universe, scientific practice, mathematical ontology, formal epistemology, Gödel phenomena

A collaboration with GPT-5.1-Thinking-Extended. 51 pages. CC4.0.

Outline

1. Introduction: A Civilization Inside Peano Arithmetic
1.1. Motivating the thought experiment
1.2. Peano arithmetic as a foundational ceiling rather than a starting point
1.3. Scope and structure of the paper

2. What It Means to Be "Limited to Peano Arithmetic"
2.1. Ontological restriction: numbers and encoded structures
2.2. Proof-theoretic restriction: only PA as the official deductive system
2.3. Definability and expressiveness in the language of arithmetic
2.4. Comparing PA to stronger foundational systems

3. Computation and Algorithmic Reasoning Within PA
3.1. Arithmetization of Turing machines and programs
3.2. Formalizing effective procedures and basic complexity theory
3.3. What PA can prove about algorithms in everyday use
3.4. Undecidability and the halting problem viewed from inside PA

4. Ordinary Mathematics in a PA-Limited World
4.1. Coding integers, rationals, and reals arithmetically
4.2. Real analysis and calculus formalizable in PA
4.3. Finite combinatorics, graph theory, and algebra over N
4.4. Reverse mathematics and the surprising reach of weak systems

5. Physics and Engineering Under Arithmetic Foundations
5.1. Modeling continuous systems via discrete and countable encodings
5.2. Differential equations, stability, and control theory in arithmetic clothes
5.3. Quantum mechanics on separable spaces and effective approximations
5.4. Numerical simulation as a paradigmatic PA-friendly practice

6. The Invisible Ceilings: True but Unprovable Statements
6.1. Gödel incompleteness as a structural feature of PA
6.2. Examples from combinatorics and number theory
6.3. Program termination, safety proofs, and verification gaps
6.4. The boundary between practical undecidability and foundational incompleteness

7. Higher Set Theory and the Mathematics That Goes Missing
7.1. Measure, compactness, and choice beyond PA's strength
7.2. Functional analysis and structures requiring stronger comprehension
7.3. Large cardinals, forcing, and independence phenomena
7.4. How little of this is needed for current empirical science

8. Epistemic Life Inside PA: Knowledge, Certainty, and Blind Spots
8.1. What counts as a proof in a PA-bounded culture
8.2. The impossibility of proving the consistency of one's own foundations
8.3. Informal knowledge versus formal certifiability
8.4. Scientific realism and mathematical Platonism under PA constraints

9. Comparing PA-World Science with Our Own
9.1. Would experimental and technological trajectories differ?
9.2. Where frontier pure mathematics diverges
9.3. Possible long-term effects on education and conceptual imagination
9.4. Hybrid strategies: unofficial stronger reasoning, official PA bookkeeping

10. Broader Philosophical Implications
10.1. Limits of formalization as a feature, not a bug
10.2. Arithmetic as a window onto but not a capture of mathematical truth
10.3. The meaning of "completeness" for a scientific civilization
10.4. Lessons for current debates on foundations and mechanized reasoning

11. Conclusion
11.1. How much can be accomplished inside PA
11.2. Where the arithmetic ceiling actually matters
11.3. Open questions for logic, philosophy, and the practice of science
11.4. Final reflections on living with necessary incompleteness

12. References

1. Introduction: A Civilization Inside Peano Arithmetic

A natural way to think about mathematical foundations is as a ladder: we start with something simple like Peano arithmetic, then climb upward to richer systems, eventually reaching the towering heights of set theory, category theory, large cardinals, and beyond. In practice, that is how much of twentieth-century logic and foundations has been narrated. Peano arithmetic appears as a low rung, useful but obviously incomplete, something to be transcended rather than inhabited.

This paper asks a different question. Instead of treating Peano arithmetic as a primitive starting point, suppose it were a ceiling. Imagine a scientific and technological civilization whose official mathematics is permanently bounded by the axioms and rules of first-order Peano arithmetic. Every formal proof, every certified theorem, every piece of "hard" justification is required to fit into that framework. Anything that cannot be encoded and proved within Peano arithmetic may be discussed informally, but it never attains the status of officially recognized knowledge.

The thought experiment is not about crippling intelligence or banning creativity. It is about examining how far a modest formal system really reaches, and where its limits begin to have concrete implications for science, technology, and philosophical self-understanding. As contemporary logic has revealed, Peano arithmetic stretches surprisingly far into what we normally regard as advanced mathematics. At the same time, incompleteness and proof-theoretic boundaries guarantee that certain truths and certain kinds of reasoning always remain beyond its grasp. Living inside Peano arithmetic is therefore neither trivial nor catastrophic, but it is distinctive. The goal of this paper is to understand that distinctiveness.

1.1. Motivating the thought experiment

There are at least three motivations for taking this imaginary PA-limited civilization seriously.

First, it exposes how much of modern science and engineering rests on finitary, arithmetizable reasoning. Claims about algorithms, numerical methods, finite combinatorial structures, and effective approximations can be expressed in the language of arithmetic. Formalizing them within Peano axioms forces us to see to what extent our current practices are already arithmetical in spirit, whether we acknowledge it or not.

Second, it sharpens our understanding of incompleteness. Gödel's theorems and later developments often appear as abstract limitations, disconnected from everyday scientific life. By connecting them to questions about program verification, long-term mathematical research, and the formal status of certain combinatorial principles, the thought experiment makes these limitations concrete. What does it mean for a scientific community if some arithmetical truths about its own algorithms and models are, in principle, unprovable within its foundational system?

Third, the scenario illuminates the gap between what a civilization can actually do and what it can formally justify. A PA-limited world might still build sophisticated computers, design quantum devices, and explore the cosmos. The restriction is not to intelligence or experimentation but to formal certification. That gap between practice and provability is already present in our world. Treating Peano arithmetic as a binding ceiling lets us study it in a controlled way, without distractions from still stronger systems.

1.2. Peano arithmetic as a foundational ceiling rather than a starting point

In most foundational narratives, Peano arithmetic is introduced early and then quickly surpassed. Stronger axiomatic systems are brought in to capture real analysis, measure theory, functional analysis, or set theory. The impression is that serious mathematics, and therefore serious science, needs more than Peano axioms can offer.

The PA-limited world flips this picture. Here, Peano arithmetic is not a primitive stepping stone but the ultimate court of appeal. One can still speak informally about sets of real numbers or about infinite dimensional spaces, but such talk acquires official force only when translated into and justified within arithmetic. Any appeal to stronger axioms is, from the point of view of the civilization's formal culture, a kind of metaphysical speculation.

Treating Peano arithmetic as a ceiling raises subtle questions. How much of real analysis can be reconstructed arithmetically? Which standard results in physics depend, in their proofs, on principles that outrun PA? How often do methods from stronger systems enter scientific reasoning, even if implicitly? The ceiling metaphor emphasizes that the restriction is not merely linguistic. It binds what counts as proved, what can be taught as fully established, and what scientists can rely on when they demand formal guarantees.

At the same time, this perspective avoids caricaturing PA as a cramped or primitive environment. By now it is well understood that large regions of "ordinary" mathematics, especially when restricted to countable and effectively described objects, can be developed in systems no stronger than Peano arithmetic or only modestly stronger. The ceiling is higher than it first appears, which makes its eventual height and its limitations all the more interesting.

1.3. Scope and structure of the paper

The remainder of the paper explores the PA-limited civilization along three intertwined axes: technical capability, mathematical reach, and epistemic life.

Section 2 clarifies what it means to be limited to Peano arithmetic. It distinguishes ontological, proof-theoretic, and definability constraints, and compares PA to stronger foundational systems often assumed in contemporary mathematics.

Section 3 examines computation and algorithmic reasoning inside PA. It shows how Turing machines, effective procedures, and basic complexity theory can be formalized arithmetically, and how undecidability already appears within this framework.

Section 4 turns to ordinary mathematics. It discusses the coding of integers, rationals, and reals, the portion of analysis and combinatorics accessible within PA, and the insights from reverse mathematics that reveal just how far weak systems extend into familiar theorems.

Section 5 focuses on physics and engineering. It investigates how continuous models can be represented by discrete encodings, which differential equations and stability results are formalizable, and why numerical simulation is an archetypal PA-friendly practice.

Section 6 introduces the invisible ceilings: true but unprovable statements in arithmetic, limitations on program termination proofs, and the subtle boundary between practical undecidability and foundational incompleteness.

Section 7 reviews mathematical territories that do outrun PA, including parts of measure theory, functional analysis, and higher set theory, and asks how essential they are to empirical science.

Section 8 considers the epistemic life of a PA-bounded culture: what counts as proof, how consistency is viewed, and how informal mathematical and scientific knowledge relate to formal certifiability.

Section 9 compares PA-world science with our own, assessing likely similarities and divergences in technological trajectories and conceptual imagination. Section 10 explores broader philosophical implications for our understanding of formalization, completeness, and the role of arithmetic in scientific thought. Section 11 concludes by summarizing how much can be accomplished inside PA and where its status as a ceiling becomes decisive.

2. What It Means to Be “Limited to Peano Arithmetic”

To say that a civilization is limited to Peano arithmetic can mean several different but related things. It can be a statement about what kinds of objects the civilization officially recognizes as existing, about what kinds of proofs it accepts as valid, or about what kinds of statements it can even formulate in its official language. These are ontological, proof-theoretic, and definability constraints respectively.

In practice, a PA-limited world would mix all three. Its textbooks, formal proofs, automated proof assistants, and official mathematical culture would all be built around a single, fixed framework: first-order arithmetic with equality, zero, successor, addition, multiplication, and order, together with the Peano axioms and the induction schema. Everything else, from real

numbers to spaces of functions, would be present only insofar as it can be encoded arithmetically.

This section unpacks the three main senses of limitation and then situates Peano arithmetic relative to stronger foundational systems. Doing so clarifies both how generous PA actually is and in what sense it really does act as a ceiling.

2.1. Ontological restriction: numbers and encoded structures

An ontological restriction is a claim about what sorts of entities are allowed to exist in the official mathematical universe. In a PA-limited world, the basic objects are just the natural numbers 0, 1, 2, and so on. There are no primitive sets, functions, or infinite sequences. If one wants to speak of those, one must do so indirectly by encoding them as numbers.

This encoding practice is familiar from ordinary logic and computability theory. A finite string of symbols can be coded as a single integer, for example via its prime factorization or via a positional numeral system with a large base. A pair (a, b) can be coded as a single number using a pairing function. Finite graphs, automata, formulas, proofs, programs, and even finite pieces of analysis can all be represented as integers via systematic schemes.

In a PA-limited civilization, this coding is not a convenience but a necessity. To define a real number, one does not introduce a new primitive kind of object; one introduces a predicate on the natural numbers that picks out which numbers code Cauchy sequences of rationals with a desired convergence property. A function from N to N becomes a set of ordered pairs, which in turn becomes a set of codes for those pairs. A probability distribution is an assignment of rational weights to finitely or countably many codes representing events.

From the inside, mathematicians in such a world would become adept at thinking of higher structures as patterns in the arithmetic of codes. A topological space is a family of subsets of N satisfying certain axioms, each subset presented by a characteristic function that is itself a coded object. A differential equation is a relation among codes for functions satisfying discrete approximations of the derivative. The ontology never expands beyond the natural numbers, but the patterns definable over them become the stand-ins for richer worlds.

The ontological restriction therefore does not prohibit richer mathematics. Instead it forces a discipline: nothing exists unless it can be faithfully represented arithmetically. The cost is conceptual weight at the foundational level; the benefit is that the entire mathematical universe is grounded in a single, robust domain.

2.2. Proof-theoretic restriction: only PA as the official deductive system

The second dimension of limitation is proof-theoretic. Here the claim is not just that all objects are coded into N, but that all officially accepted proofs must be derivable from the axioms and rules of Peano arithmetic, formulated in its first-order language.

This has several consequences.

First, any chain of reasoning that appeals to principles not derivable in PA is treated as informal or heuristic. Mathematicians might still find it persuasive to argue using, say, the full axiom of choice or transfinite induction up to certain ordinals, but as far as the official record is concerned, these arguments do not establish theorems until they are recast in purely arithmetical terms and their logical steps are verified to be PA-derivable.

Second, the proof-theoretic restriction sharply delimits what can be proved about arithmetic itself. Gödel's incompleteness theorems apply to PA, so there are true arithmetical sentences that are not provable within it. In particular, PA cannot prove its own consistency. A PA-limited civilization therefore has a built-in, non-removable limit on how far it can go in formally assuring itself that its foundational system is sound.

Third, proof-theoretic limitation constrains formal verification of algorithms and systems. Many interesting properties of programs, especially those that amount to universal statements over all inputs (for all n, the program halts or behaves safely), are of the same logical form as the incompleteness examples. Some of these properties may be true in the standard model of arithmetic but unprovable in PA. In a civilization whose official deductive power is bounded by PA, such properties can never receive fully formal certification.

Despite these limitations, PA is far from weak. It can prove the correctness of many concrete algorithms, establish induction for a wide range of properties, and support complex combinatorial arguments. The proof-theoretic restriction therefore does not wipe out the possibility of rigorous reasoning. It does, however, ensure that there are permanently unreachable truths even in the narrow domain of the natural numbers.

2.3. Definability and expressiveness in the language of arithmetic

The third form of limitation concerns definability: what can be said at all in the official language. Peano arithmetic is formulated in a fairly sparse first-order language, typically with symbols for 0, the successor function S, addition, multiplication, equality, and sometimes an order relation. There are variables for numbers but not for sets of numbers or functions from numbers to numbers. Quantifiers range only over individual natural numbers.

This language is remarkably expressive when combined with arithmetic coding. One can write formulas that say things like "the number n encodes a sequence that is a proof of the formula

with code m" or "for every input x, the computation of program p on input x eventually halts with output y." Through these encodings, PA can internalize talk about syntax, computation, and even its own deductive activity.

However, the first-order nature of the language limits the kinds of predicates that can be directly defined. Properties that essentially quantify over all subsets of N or all functions from N to N cannot be captured in a single formula without further coding tricks. Statements about uncountable sets or about arbitrary real functions must be paraphrased into statements about codes and finite or countable approximations.

Definability has both technical and philosophical consequences in a PA-limited world. Technically, it shapes which mathematical concepts can be directly formulated and which require multi-layered encodings. Philosophically, it affects how mathematicians think about their objects. Real numbers become certain equivalence classes of Cauchy sequences of rationals; these sequences are themselves given by definable rules or by sets of pairs of integers. Higher structures lose some of their immediate geometric or set-theoretic flavor and are instead perceived as complicated but precise patterns in the arithmetic of codes.

The expressiveness of the language of arithmetic thus acts as a kind of lens. It does not prevent mathematicians from accessing higher notions, but it forces all such access to pass through arithmetical descriptions. The world of mathematics is seen, always, through the medium of N and its definable relations.

2.4. Comparing PA to stronger foundational systems

To understand what it means to be limited to PA, it is useful to contrast it with stronger systems that play central roles in contemporary foundations.

One benchmark is second-order arithmetic. Here, in addition to variables for individual numbers, one has variables for sets of numbers and can quantify over them. The axioms assert that the natural numbers satisfy the Peano properties and impose comprehension principles that guarantee the existence of sets defined by certain formulas. Subsystems of second-order arithmetic form a hierarchy of strength, with some only slightly stronger than PA and others comparable to substantial portions of set theory.

Another benchmark is full set theory, such as Zermelo–Fraenkel with choice. In these systems, the universe consists of sets built up in transfinite stages, and the natural numbers are just one small definable part of a vastly richer ontology. Real numbers are sets; functions are sets; topological spaces, vector spaces, and measures are all set-theoretic constructions. The axioms of set theory assert the existence of infinite hierarchies of sets with specific closure properties, and additional axioms can postulate large cardinals with strong combinatorial features.

From a proof-theoretic point of view, these stronger systems can prove many statements about natural numbers that PA cannot. For example, certain combinatorial principles and termination facts for specific classes of programs are provable in relatively modest extensions of PA but not in PA itself. From an ontological point of view, they directly posit entities, such as arbitrary subsets of R or large cardinals, that have no straightforward arithmetical coding.

In a PA-limited world, the role of these stronger systems would be ambiguous. They might be studied as heuristic devices or as external meta-theories used to reason about what is true in arithmetic, but the results they yield would only be accepted as official theorems when they can be translated into and verified by PA. The stronger systems would resemble powerful but unofficial instruments, akin to speculative physical models that guide experiment but are not yet accepted as part of the established theory.

The comparison highlights a central tension. On the one hand, PA is strong enough to capture a surprising amount of the mathematics actually used in science and engineering. On the other hand, stronger foundational systems enable proofs and concepts that PA cannot reproduce. Treating PA as a foundational ceiling means accepting that some of those proofs and concepts, however compelling they may seem, will remain forever outside the sphere of formally sanctioned knowledge.

3. Computation and Algorithmic Reasoning Within PA

If there is any area of modern thought that seems tailor-made for Peano arithmetic, it is the theory of computation. Algorithms are finite strings of symbols operating on finite strings of symbols. They manipulate integers, lists, graphs, and other discrete structures that are easily coded by natural numbers. The basic claims we make about them are arithmetical in form: for every input n, the program halts; the output is always less than 2n; no two inputs produce the same hash.

In a PA-limited civilization, computation is therefore not a secondary application of arithmetic but one of its central expressive domains. The natural numbers provide the universe of codes; the axioms of PA provide the rules by which one reasons about execution steps, resource bounds, and correctness. This section sketches how Turing machines and programs are arithmetized, how effective procedures and simple complexity notions are formalized, what PA can prove about everyday algorithms, and how undecidability and the halting problem look from the inside.

3.1. Arithmetization of Turing machines and programs

The basic trick of arithmetization is to replace every syntactic or combinatorial object with a unique natural number, and every relation among such objects with a numerical relation among their codes. In the context of computation, the key objects are:

- machine descriptions (or program texts)
- machine states and tape configurations
- finite execution histories or single computation steps

Each of these can be encoded as a number in many different but equally effective ways. One standard pattern is:

1. Fix a finite alphabet of symbols, including tape symbols, state identifiers, and instruction opcodes. Assign each symbol a small positive integer.
2. Encode a finite sequence of symbols s0, s1, …, sk as a single number using, for example, a positional encoding
   n = s0 + s1 * B + s2 * B^2 + … + sk * B^k
   where B is larger than any symbol code.
3. Encodings for two-dimensional or more complex objects, such as configurations of a Turing machine with a marked head position, can be constructed from encodings of simpler pieces using pairing functions or nested positional schemes.

Inside PA, these encodings are just functions from N to N defined by primitive arithmetic formulas. One can write down a formula Step(m, c, c′) that says: “c′ is the configuration obtained from configuration c by applying one transition rule of the machine with code m.” This relation is primitive recursive; PA can prove basic properties about it, such as that for each machine and configuration there is at most one next configuration.

By iterating Step, one obtains a relation Reach(m, c0, t, c) that says: “starting from initial configuration c0 of machine m, after exactly t steps the machine is in configuration c.” This relation is also definable by an arithmetic formula. The halting relation Halt(m, x, t) that says “machine m on input x halts exactly at time t” can then be written in terms of Reach and a predicate that recognizes halting configurations.

From the point of view of a PA-limited civilization, this means that all talk of “running a program for t steps” or “reaching a halting state” can be translated into statements about natural numbers satisfying certain formulas. Turing machines and general programs are not exotic or external; they are particular families of coded objects whose behavior is governed by arithmetic relations. The machinery of computability theory lives entirely inside the arithmetic universe.

3.2. Formalizing effective procedures and basic complexity theory

Once programs are arithmetized, the notion of an effective procedure can be captured in several equivalent ways. A function f: N -> N is called computable if there is a program p such that for all inputs n, the execution of p on input n halts with output f(n). Formally, this can be expressed inside PA with a formula:

For all n there exists t and y such that Halt(p, n, t) and Output(p, n, t, y),

where Output expresses that the configuration at time t encodes output y. The condition that the output is unique, and that the computation halts for every input, can also be rendered by arithmetic formulas.

Different models of computation, such as register machines, lambda calculus, or recursive function schemata, can all be arithmetically formalized and shown equivalent within PA. For instance, the class of primitive recursive functions can be defined inductively from basic initial functions using composition and primitive recursion, and PA can prove that every such function is computable by some program in the above sense. Conversely, standard proofs that every computable function is recursive can be carried out in arithmetic.

Basic complexity notions can also be imported. Given a program p, one can define its time complexity function T_p(n) as the least t such that p halts on input n in t steps, if such a t exists. This is a function from N to N or partially defined on some inputs; PA can express such definitions using bounded quantifiers and minimization operators coded arithmetically. Statements like “T_p(n) is bounded by a polynomial in n” become arithmetic sentences:

There exist constants c and k such that for all n and t, if Halt(p, n, t) then t <= c * n^k.

Similar translations work for space complexity, number of states visited, or other combinatorial resource measures.

What PA cannot do is decide in general, for an arbitrary program p and resource bound B(n), whether T_p(n) <= B(n) for all n. That is a deep problem, related to the undecidability of many global complexity questions. But it can faithfully represent the definitions and prove numerous concrete theorems about specific algorithms and classes of algorithms.

In a PA-limited civilization, then, basic complexity theory is a perfectly legitimate branch of arithmetic. The formal statements about machines, time bounds, and resource trade-offs are all sentences about N, subject to the usual Peano axioms and induction. The notion of what is effectively computable cannot be formalized more precisely than this, even in stronger systems, so at this level PA is not obviously deficient.

3.3. What PA can prove about algorithms in everyday use

The ultimate test for a foundational system is not whether it can represent definitions, but whether it can support the correctness proofs and reasoning patterns that practitioners care about. For computation, this mainly concerns termination and partial correctness.

Consider familiar algorithms:

- Euclid's algorithm for greatest common divisors
- standard sorting algorithms like insertion sort or merge sort
- basic graph search such as depth-first search or breadth-first search
- elementary number-theoretic algorithms such as trial division, the Sieve of Eratosthenes, or simple primality tests
- linear algebra routines over finite matrices, such as Gaussian elimination

Each of these algorithms manipulates finite discrete structures in a way that can be coded arithmetically. Their correctness proofs typically proceed by induction on the size of the input, by loop invariants, and by straightforward combinatorial reasoning. All of these proof techniques can be formalized within PA.

For Euclid's algorithm, for instance, PA can prove that the sequence of remainders strictly decreases and therefore must terminate, and that the final nonzero remainder divides both inputs and is divisible by any other common divisor. This uses only basic properties of division and induction on the size of the remainder.

For sorting algorithms, PA can express the key invariants: that the multiset of elements in the array is preserved, and that the array is nondecreasing at termination. The induction principles required are over finite sequences, which can be encoded as numbers, and the reasoning needed is entirely within the reach of Peano induction.

For many numerical algorithms, such as explicit finite-difference schemes for differential equations or simple iterative solvers, PA can formalize the proofs that each step is defined, that arrays remain within their index bounds, and that the procedure terminates after a fixed number of iterations. Where convergence as step size goes to zero is at issue, stronger analytical tools enter the picture, but at the level of any concrete, finite computation, the reasoning is still arithmetical.

Even in more advanced domains like cryptography and formal verification, much of the reasoning about concrete constructions lives at the level of PA. Arguments that a particular block cipher has no fixed points with certain properties, or that a protocol prevents a specific

class of simple attacks, can be unfolded into finite combinatorial cases and probabilistic bounds over finite spaces. As long as the properties in question are finitary, and the proofs proceed by concrete combinatorial and number-theoretic reasoning, PA can accommodate them.

In practice, this means that a PA-limited civilization retains the ability to formally verify substantial portions of its everyday software and hardware infrastructure, as long as specifications themselves are arithmetically expressible. The limitation is not that most algorithms become unverifiable, but that some global properties of very general algorithms and systems, especially those involving unbounded behavior over all inputs, will escape formal proof.

3.4. Undecidability and the halting problem viewed from inside PA

The halting problem is often presented as a limit on what any one algorithm can do: there is no single computable procedure that decides, for every program p and input n, whether p halts on n. From the perspective of PA, this undecidability result is itself an arithmetical theorem. The formalized argument that no such total deciding program exists can be carried out within arithmetic, using the arithmetized notion of computation and a diagonal construction.

Inside a PA-limited civilization, therefore, logicians can prove that no program H has the property:

For all p and n, H(p, n) halts and outputs 1 if p halts on n, and 0 otherwise.

The argument shows that assuming the existence of such a program leads to a contradiction that PA can detect. This is a powerful internal result: the community knows, at the level of its official foundation, that its own computational methods cannot solve certain global decision problems.

At the same time, the halting problem has a second, subtler face when viewed through incompleteness. For any particular program p, the statement

For all n, p halts on input n

is an arithmetical sentence of a specific logical form. Some such sentences are provable in PA, some are refutable, and some are independent: true in the standard model of N but unprovable in PA. The existence of these independent termination statements is a consequence of Gödel's incompleteness theorems.

From inside PA, one can formulate these statements, reason about them, and sometimes derive conditional results of the form: "If PA is consistent, then program p does not halt on any input" or "If PA is consistent, then PA cannot prove that p always halts." What PA cannot do is cross

certain boundaries: it cannot prove its own consistency, and it cannot settle all such termination statements.

For a PA-limited civilization, the halting problem thus appears in two layers. At the global layer, it is a theorem: no universal halting decider exists, and this negative result is itself provable in the official system. At the particular layer, the community confronts specific programs whose termination behavior is true-but-unprovable or false-but-unprovable in PA. Some of these programs may be artificially contrived to encode logical paradoxes, but others can arise naturally from combinatorial constructions.

This double aspect has practical implications. In the design of critical systems, engineers may seek formal guarantees that certain processes always terminate or remain within safe bounds. For many such systems, PA can supply the needed proofs. For some, especially those whose behavior effectively encodes complex mathematical principles, no such proof will ever exist inside PA, even if the systems are in fact safe. The foundational ceiling therefore translates into an unavoidable gray zone in formal verification: a region where the truth about algorithmic behavior outruns what can be certified.

In summary, computation gives a vivid illustration of life inside Peano arithmetic. On the one hand, the entire theory of effective procedures, from Turing machines up through substantial parts of complexity theory, can be housed within PA. On the other hand, undecidability and incompleteness ensure that there are algorithms whose global behavior cannot be settled, even in principle, by the civilization's official mathematics. The arithmetic universe is rich enough to contain both the power of computation and its own structural limits.

4. Ordinary Mathematics in a PA-Limited World

So far the focus has been on computation, where the fit with Peano arithmetic is almost immediate. Ordinary mathematics might seem less obviously arithmetical. Real numbers, continuous functions, differential equations, and geometric objects are usually introduced through pictures and set-theoretic language, not through encodings into the natural numbers.

A PA-limited civilization would still teach and use these concepts in their familiar informal forms. Analysts would still talk about intervals, limits, integrals, and Hilbert spaces. What changes is the way those concepts are grounded when one asks for formal definitions and proofs. The question becomes: which parts of everyday mathematics can be translated into arithmetic and verified within PA, and which require stronger principles?

The answer is again two-sided. On the one hand, an impressive amount of real analysis, combinatorics, and algebra can be formalized in arithmetical systems no stronger than PA, or

only modestly stronger. On the other hand, certain compactness principles, existence theorems, and infinitary constructions lie beyond what PA can prove. This section surveys both sides, starting with the basic encodings of number systems and moving outward to analysis, combinatorics, and the landscape revealed by reverse mathematics.

4.1. Coding integers, rationals, and reals arithmetically

Within the natural numbers, the integers, rationals, and reals are not new primitive domains but constructions.

Integers can be represented, for example, as ordered pairs of natural numbers (a, b) interpreted as a − b, modulo the equivalence relation that identifies (a, b) and (c, d) when a + d = b + c. The operations of addition and multiplication are defined on these pairs in the usual way and provably respect the equivalence relation. All of this can be spelled out in arithmetic, using a fixed pairing function to code each pair as a single natural number.

Rational numbers can be coded as pairs of integers (p, q) with q nonzero, again modulo the equivalence relation p1 q2 = p2 q1. Positive rationals can be represented by pairs of natural numbers, and negative rationals by signed pairs or by embedding integers into the same scheme. Operations are given by the standard fraction rules. PA can prove that these operations are well defined on equivalence classes and satisfy the field axioms.

Real numbers are more delicate, but they too can be brought inside arithmetic. One common approach is to define a real as an equivalence class of Cauchy sequences of rationals with a specified modulus of convergence. Each Cauchy sequence is a function from N to Q, which can be coded by a single natural number describing an algorithm to compute its terms. Two sequences are equivalent if their difference converges to zero, a condition expressible by an arithmetical formula quantifying over indices and rational bounds.

Another approach is via Dedekind cuts in the rationals, represented as sets of rationals (or codes for rationals) satisfying certain properties. Since PA has no primitive sets, such a cut is represented by a unary predicate on the codes of rationals, or by a number coding that predicate. Again, the defining properties of a cut can be captured arithmetically.

For a PA-limited civilization, these constructions are not exotic exercises but the official backbone of their analysis. When a textbook says “let x be a real number,” the formal counterpart is “let x be a code for an equivalence class of suitably convergent rational sequences,” or “let x be a code for a Dedekind cut.” Limits, continuity, and differentiability become statements about these codes, written out in the language of arithmetic.

This does not mean that mathematicians must constantly think about codes. In practice, they reason with reals in the usual conceptual way, appealing to the arithmetical encodings only

when they need to certify that a definition or proof can be made precise in PA. The ontology is still purely numerical, but the conceptual landscape of analysis is reconstructed within it.

4.2. Real analysis and calculus formalizable in PA

Given these encodings, how much of ordinary analysis can be developed inside Peano arithmetic? Surprisingly, quite a lot.

Core notions of calculus, such as limits of sequences, continuity of functions on intervals, differentiability, and the Riemann integral, are all definable arithmetically. A sequence of reals is a function from N to coded reals; convergence is expressed in terms of quantifiers over natural indices and rational error bounds. A function f from reals to reals can be represented by an algorithm that, given a code for a real x and a precision parameter, produces a code for an approximation to f(x) at that precision. Continuity at a point becomes the familiar epsilon–delta condition, translated into statements about the outputs of such algorithms.

Within this framework, PA can prove many standard theorems of elementary analysis. The limit laws for sums and products of convergent sequences are straightforward. The fundamental results about continuity, such as the fact that continuous functions on closed intervals are bounded and attain their bounds, can be proved when those intervals and functions are presented in effective or arithmetically definable ways. The intermediate value theorem, in constructive or semi-constructive formulations, can be established for functions that are given by explicit algorithms and meet suitable regularity conditions.

Differentiation and integration can also be handled when restricted to functions that are defined by arithmetically describable procedures. For example, the mean value theorem and the fundamental theorem of calculus hold for many classes of functions encountered in applied analysis, provided the functions are represented in a way that allows their derivatives or integrals to be approximated algorithmically. Existence and uniqueness theorems for solutions to certain differential equations, under Lipschitz and boundedness conditions, can be proved in relatively weak systems, and their proofs can be translated into PA.

What tends to fall outside PA, or require stronger axioms, are fully general statements that quantify over all bounded sequences of reals, all continuous functions on a compact interval, or all measurable sets. Strong forms of compactness, such as the Heine–Borel theorem for arbitrary covers or the Bolzano–Weierstrass theorem for all bounded sequences, may require comprehension principles that exceed what PA can justify. But as long as analysis is tethered to explicitly given, effectively described objects, the core theorems underpinning numerical methods and much of mathematical physics can live in the arithmetic setting.

From the standpoint of physics and engineering, this is crucial. The mathematical tools used to model and analyze real systems typically involve specific differential equations, specific function

spaces, and specific approximation schemes. These can be represented in PA and reasoned about there. The loss concerns abstract uniform principles that range over all possible objects of a certain type, not the concrete theorems most central to practice.

4.3. Finite combinatorics, graph theory, and algebra over N

Finite combinatorics and graph theory are almost tailor-made for a PA-limited world. Finite sets and finite graphs can be encoded as numbers, and properties such as connectivity, colorability, or existence of particular substructures become arithmetical predicates. For example, a finite graph on n vertices can be represented by an n-by-n adjacency matrix, which can in turn be encoded as a binary string and then as a natural number.

Statements such as “every finite graph with property P also has property Q” can therefore be interpreted as universally quantified sentences over N. Many classical results in graph theory and combinatorics, including bounds on Ramsey numbers for small parameters, Hall’s marriage theorem in finite settings, or existence of matchings and coverings in appropriately structured graphs, can be proved using induction and combinatorial arguments that PA can formalize.

Similarly, algebraic structures whose underlying sets are finite or finitely generated and whose operations are effectively given can be housed in PA. Finite groups, rings, and fields can be represented by multiplication tables or generating sets together with relations. Theorems such as Lagrange’s theorem for finite groups, classification of finite abelian groups, or basic results on finite fields are provable in relatively weak systems and can be arithmetized.

Even for infinite algebraic structures, much of the reasoning can be grounded in finitary approximations. For instance, many structural theorems about countable groups or rings can be rephrased in terms of properties of finite generating sets and relations among them. PA can then express and prove these properties as arithmetical statements.

There are, however, boundaries. Certain combinatorial principles that look purely finitary turn out to lie beyond PA. Strengthened versions of Ramsey-type theorems, which assert the existence of large homogeneous sets under certain colorings, are known to be true in the standard natural numbers but unprovable in Peano arithmetic. These results reveal that even in the realm of finite combinatorics, the PA ceiling can be reached.

From the internal perspective of a PA-limited civilization, this would show up as follows: combinatorists can derive and verify many theorems, and for small cases they can check instances by computation. At some point, they encounter principles whose general validity they strongly suspect but cannot prove using the accepted axioms. Stronger set-theoretic or second-order methods might suggest that these principles hold, but any attempt to translate those methods back into PA runs into incompleteness. The field of finite combinatorics thus becomes a site where the foundational ceiling is felt directly.

4.4. Reverse mathematics and the surprising reach of weak systems

Reverse mathematics is a program in the foundations of mathematics that systematically analyzes which axioms are needed to prove which theorems, typically working over subsystems of second-order arithmetic. Although its formal setting lies slightly above PA, its results shed light on what can be accomplished in PA or systems very close to it.

The basic pattern in reverse mathematics is to fix a weak base theory that captures a minimal amount of arithmetic and then show that many familiar theorems of analysis, algebra, and combinatorics are equivalent, over that base, to particular additional axioms or schemes. For example, certain forms of the Bolzano–Weierstrass theorem, the Arzelà–Ascoli theorem, or basic compactness results in functional analysis turn out to be equivalent to specific comprehension principles or induction schemas.

This has two important implications for a PA-limited world.

First, many theorems that might naïvely appear to require full set theory actually only need very modest extensions of basic arithmetic. When those extensions are conservative over PA for arithmetic sentences, it means that, as far as statements purely about natural numbers go, PA already suffices. A civilization that treats PA as its ceiling can therefore adopt a great deal of "ordinary" mathematics as an informal guide, knowing that its arithmetical consequences do not outrun the foundational system.

Second, reverse mathematics highlights precisely where PA and closely related systems break down. Theorems that are equivalent to stronger subsystems, especially those involving robust compactness or comprehension over sets of reals, indicate regions where a PA-limited civilization cannot, within its official framework, recover the full generality of classical results. It might be able to prove special cases or effectivized versions, but not the universal statements about all sequences, all functions, or all covers.

The overall picture that emerges is one of surprising reach combined with sharply delineated limits. Weak systems close to PA suffice for a wide swath of the mathematics used in analysis and applications. At the same time, there are natural, robust theorems that mark the boundary beyond which PA cannot go. For a civilization living under the Peano ceiling, reverse mathematics would serve as a kind of map of its mathematical territory: a chart indicating which hills lie well within its domain, which ridges are at the frontier, and where the landscape drops off into regions that, however appealing, remain formally inaccessible.

In such a world, ordinary mathematics is not reduced to arithmetic exercises. Rather, arithmetic becomes the substrate through which analysis, combinatorics, and algebra are reconstructed. The conceptual experience of continuous and combinatorial structures remains, but their formal backing is always traceable to patterns in N.

5. Physics and Engineering Under Arithmetic Foundations

If Peano arithmetic can house much of ordinary mathematics, it is natural to ask what happens to physics and engineering, disciplines whose models are soaked in notions of continuity, differential equations, and infinite-dimensional spaces. At first glance, these seem to demand a set-theoretic universe far richer than the natural numbers. Fields and manifolds, continua and spectra, appear distant from discrete arithmetic.

The key observation is that almost all practical physics and engineering ultimately proceed by way of discrete approximations. Differential equations are solved numerically on grids, continuous signals are sampled and digitized, and operators on infinite-dimensional spaces are approximated by large but finite matrices. Measurements themselves are finite strings of digits from instruments. Beneath the smooth equations lies a lattice of integers and rational approximations.

In a PA-limited civilization, this lattice is the foundational bedrock. The official mathematics of physics and engineering is phrased in terms of codes for discrete and countable structures, together with theorems about how well these approximate idealized continuous models. The language of continua survives as a guide and shorthand, but when one asks what is really being proved, the answer is always arithmetic.

This section explores how continuous systems are modeled via discrete encodings, how differential equations and control theory look in arithmetic clothes, how quantum mechanics on separable spaces fits into this framework, and why numerical simulation is an almost perfect example of a PA-friendly practice.

5.1. Modeling continuous systems via discrete and countable encodings

A central task in physics is to describe how physical quantities vary in space and time. Classical field theories assign real-valued functions to points in a continuum; fluid dynamics describes densities and velocities as functions on regions of space; electromagnetism uses vector fields and potentials. At the level of idealized equations, these objects are infinite collections of real numbers indexed by a continuum of points.

To bring such models into arithmetic, two steps are needed: discretization and coding.

Discretization replaces the continuum by a countable or finite grid. Instead of a function defined on every real x in an interval, one considers its values on rational points or on a mesh of points with spacing h. In space, one approximates a region by a lattice of volumes; in time, one breaks evolution into discrete steps of size Δt. This is standard practice in numerical analysis and computational physics.

Coding then replaces each discrete structure by a natural number that encodes its data. A function on a finite grid can be represented by a list of rational values, which in turn is coded as a sequence of integers. A time-evolving field on a lattice becomes a sequence of such coded configurations. Boundary conditions and source terms are likewise encoded.

Within PA, one does not assert the existence of all possible continuum fields. Instead one asserts the existence of algorithms that, given a desired resolution and range, produce discrete approximations satisfying certain properties. For example, one might formalize statements like:

For every integer N, there exists a coded grid of size N with associated field values such that the discrete Laplace operator applied to these values approximates the continuous Laplacian of a given function to within error 1/N.

The function in question would itself be given by an algorithm producing rational approximations. In this way, the continuous model is replaced by a family of discrete approximations with arithmetically expressible convergence behavior.

In practice, physicists in a PA-limited world would still draw continuous curves and surfaces, still write down partial differential equations in their familiar form, and still reason using analogies and geometric pictures. The underlying formalism, however, would treat these as shorthand for statements about the existence and behavior of discrete encodings at arbitrary but finite resolutions. The continuum becomes a regulative ideal: something approached in the limit but never needing to exist as a completed uncountable structure in the foundational theory.

This perspective aligns surprisingly well with how physical theories are used in our own world. Experiments always have finite precision; simulations always run on finite hardware; predictions are always expressed within error bars. PA provides a natural language for these finitistic aspects of modeling.

5.2. Differential equations, stability, and control theory in arithmetic clothes

Differential equations sit at the heart of classical physics and many engineering disciplines. They describe how systems evolve in time, how fields respond to sources, and how devices behave under control inputs. In their usual form, they involve derivatives and integrals defined on continuous domains, backed by theorems about existence, uniqueness, and stability of solutions.

In a PA-limited setting, the emphasis shifts to difference approximations and effective existence results.

Take an ordinary differential equation of the form

$x'(t) = f(t, x(t))$ with $x(t0) = x0$,

where f is a function given by an explicit formula or algorithm. To handle this in arithmetic, one considers discrete time steps t0, t0 + h, t0 + 2h, and so on, and defines approximate solutions by schemes such as Euler's method or higher-order methods. The update rule

x_{n+1} = x_n + h f(t_n, x_n)

is purely algebraic and can be encoded as an arithmetic function on rational approximations to x_n.

Within PA, one can express and prove statements like:

For each natural number N, there exists a step size h <= 1/N and a sequence of rational vectors {x_n} defined by the chosen update rule such that for each time step up to some fixed horizon T, the discrete values x_n approximate a true solution of the differential equation to within error 1/N.

To make such statements precise, one needs an effective existence and uniqueness theorem: given Lipschitz conditions and bounds on f, one can construct approximate solutions whose errors can be estimated. Many such theorems are provable in relatively weak subsystems and can be translated into arithmetic. The proofs proceed by iteration, estimation of truncation errors, and induction on the number of steps, all of which are comfortable within PA.

Stability analysis and control theory also admit arithmetical formulations. The classical notion of Lyapunov stability, for example, can be expressed as: for every ε there exists δ such that if the initial condition is within δ of equilibrium, then the solution remains within ε for all future times. In discretized form, with effective approximations to solutions, this becomes a statement about inequalities on finite sequences indexed by discrete time steps. Lyapunov functions can be represented by algorithms that compute their values and derivatives (approximated by differences), and stability conditions become inequalities that PA can reason about.

Control systems, whether linear or nonlinear, are typically designed using finite-dimensional state-space models. Controllers are often implemented digitally as algorithms that update control signals based on sampled measurements. In a PA-limited world, this is an advantage: the control laws are themselves finite programs manipulating coded states and inputs. Stability and performance guarantees can be expressed as arithmetical properties of these coded trajectories.

For linear systems, properties like controllability, observability, and stability can be characterized in terms of matrices, rank conditions, and eigenvalues. While general eigenvalue theory over the reals pulls in more analysis, many practical conditions reduce to solvable algebra over polynomials with rational coefficients. PA can express and prove the required determinant and rank properties for the finite matrices involved in any given application.

The result is that a substantial portion of the differential equation and control theory toolkit, at least as used for concrete systems, can be rebuilt in arithmetic clothes. The main loss lies in very general statements about arbitrary systems and continua, not in the practice of modeling and controlling specific devices.

5.3. Quantum mechanics on separable spaces and effective approximations

Quantum mechanics often seems, at first glance, further removed from arithmetic foundations than classical physics. Its formalism involves Hilbert spaces, operators, spectra, and unbounded observables. Yet in most practical applications, the relevant Hilbert spaces are separable, and the operators of interest can be approximated by finite matrices in suitable bases.

A separable Hilbert space has a countable orthonormal basis. This basis can be indexed by natural numbers, and finite linear combinations of basis vectors with rational coefficients can be coded as finite sequences of rational numbers. Inner products of such finite combinations reduce to finite sums of rational products. Thus the algebra of finite linear combinations and finite-rank operators is easily brought into arithmetic.

Many quantum systems are modeled by Hamiltonians that, in a chosen basis, correspond to infinite matrices with entries given by explicit formulas. Approximations then proceed by truncating to large but finite subspaces. For instance, one might approximate a particle in a potential well by restricting to the first N basis states and computing the corresponding N-by-N matrix representation of the Hamiltonian. Eigenvalues and eigenvectors of this finite matrix can be computed numerically, yielding approximations to energy levels and stationary states.

Within PA, one can formalize the construction of these finite approximations and prove statements such as:

For each N, there exists a coded N-by-N matrix H_N approximating the Hamiltonian such that the first k eigenvalues of H_N differ from the corresponding eigenvalues of the full operator by at most 1/N.

To make such a statement rigorous, one needs spectral theory that relates the spectra of truncated operators to those of the full operator, under suitable conditions. Much of this theory is analytic, but many of its effective consequences for specific, well-behaved systems can be reduced to arithmetical bounds on approximations. As with classical analysis, the focus shifts to effective, algorithmically describable systems.

Measurement in quantum mechanics also naturally lends itself to PA-friendly formalization. Measurement outcomes are discrete events with probabilities given by squared amplitudes or trace expressions. For any given experimental setup, one can specify a finite set of possible outcomes and rational approximations to their probabilities. The probabilistic statements of the

theory then become arithmetical assertions about frequencies in long runs of trials and about expectations of coded random variables.

Quantum information theory, with its emphasis on qubits, quantum gates, and finite-dimensional systems, is perhaps the most congenial to PA foundations. Quantum circuits are finite diagrams built from a finite set of unitary gates. Their action on basis states can be represented as finite unitary matrices with entries that are algebraic or otherwise effectively describable complex numbers. Composition and tensor products of these matrices are finite algebraic operations. Statements about circuit behavior and algorithm correctness reduce to arithmetic on the real and complex numbers encoded via rationals and sequences.

A PA-limited civilization would therefore be capable of developing and exploiting much of quantum mechanics, at least in its separable, effectively presented form. As before, the constraint appears when one asks for very general nonconstructive existence results, arbitrary self-adjoint extensions, or spectral properties of operators not given by explicit prescriptions. But the parts of quantum theory that connect directly to experiment and technology largely live in the realm of effective approximations on countable bases.

5.4. Numerical simulation as a paradigmatic PA-friendly practice

If there is a single practice that best illustrates how physics and engineering fit inside Peano arithmetic, it is numerical simulation. Simulations translate continuous models into discrete computations over finite data structures. They are implemented as programs that run on digital hardware, producing approximations to physical behavior. Their entire logic is discrete and arithmetical, even when their surface interpretation is continuous.

A typical simulation process involves several stages:

1. Model formulation: starting from physical laws, one writes down equations describing the system. These may be differential equations, integral equations, or stochastic processes.

2. Discretization: the equations are approximated on a mesh or grid in space and time. Derivatives are replaced by finite differences, integrals by sums, and continuous randomness by pseudo-random sequences.

3. Algorithm selection: one chooses numerical schemes to advance the solution in time, solve linear systems, or handle nonlinearity and constraints.

4. Implementation: the algorithms are coded in a programming language and compiled to machine code that manipulates finite arrays of floating-point numbers or rational approximations.

5. Analysis and validation: one studies convergence, stability, and error bounds, often by a mixture of theoretical estimates and empirical tests.

Every one of these steps can be expressed in arithmetic terms. The discretization is a finite data structure; the algorithms are Turing machines; floating-point representations are encodings of rational numbers with bounded denominators and fixed mantissas; error bounds are inequalities among integers and rationals. The convergence analyses, when effective, are proofs that for each desired error threshold there is a choice of mesh size and time step such that the discrete solution deviates from the ideal continuous solution by less than that threshold.

In a PA-limited world, numerical simulation would not merely be an application of mathematics but its central, fully certified mode of engaging with physical theories. The formal verification of a simulation code's basic properties, such as absence of out-of-bounds access and bounded accumulation of rounding errors, would be phrased as arithmetical theorems. More sophisticated guarantees, such as global stability or monotonicity properties of the numerical scheme, would likewise be cast as PA-expressible statements.

This has a subtle but important philosophical impact. In such a civilization, the ultimate warrant for a physical prediction is not an appeal to an ideal continuum model standing outside arithmetic, but an appeal to a proven correspondence between discrete simulations and observable outcomes, together with effective error estimates. The continuum equations are valued as guides to the choice of discretizations and as sources of qualitative insight, but the mathematics that is actually proved lives in the arithmetic of codes and discrete processes.

This is already close to the situation in high-performance computing and applied engineering today, where simulations of weather, materials, plasmas, or cosmology rely entirely on discretizations and finite computations. The foundational difference is that in our world, these practices float on top of a richer set-theoretic picture, whereas in a PA-limited world they would be the fundamental, officially recognized core.

In this sense, numerical simulation is the paradigmatic PA-friendly practice because it makes explicit the route from continuous models to discrete, coded computations. It embodies the idea that what science ultimately uses are families of finite approximations under control, rather than completed infinities. A civilization under arithmetic foundations would find in numerical simulation not a compromise, but a natural expression of its deepest mathematical commitments.

6. The Invisible Ceilings: True but Unprovable Statements

Up to this point, Peano arithmetic has looked surprisingly generous. It can encode computation, reconstruct large stretches of analysis, support combinatorics and algebra, and underwrite most of the mathematics used in physics and engineering practice. If one stopped there, it might be tempting to think of PA as simply a lean but adequate foundation.

Incompleteness breaks that illusion. It says that, no matter how far one pushes within PA, there are truths about the natural numbers that the system can never reach. Some of these truths are highly contrived; others are strikingly natural combinatorial principles or innocuous-looking termination claims about specific algorithms. For a PA-limited civilization, these unreachable truths constitute an invisible ceiling: they sit above the formal system, shaping its limits, but cannot be seen from within as theorems.

This section examines incompleteness as a structural feature of PA, gives examples from combinatorics and number theory, connects them to program verification, and then distinguishes foundational unprovability from more familiar, practical forms of undecidability.

6.1. Gödel incompleteness as a structural feature of PA

Gödel's incompleteness theorems are often stated in a general way: any sufficiently strong, effectively axiomatized, consistent theory of arithmetic is incomplete. Peano arithmetic is the canonical example of such a theory. It is strong enough to formalize basic arithmetic, to talk about its own syntax, and to interpret the workings of simple computing devices. Its axioms can be listed by a mechanical procedure. Assuming consistency, Gödel's theorems apply.

The core idea is that PA can talk about its own proofs. Using arithmetization, formulas, proofs, and derivations are encoded as natural numbers, and the relation "n is the code of a valid PA-proof of the formula with code m" becomes an arithmetically definable predicate. With this machinery in place, PA can express statements about which formulas are provable and which are not.

Gödel then constructs a particular sentence that, when decoded, says something like: "This sentence has no proof in PA." If PA could prove that sentence, it would be proving a false statement about its own proof system, contradicting the assumption of consistency. If PA could prove its negation, it would be proving that there is a proof where there is none, again contradicting consistency. So, assuming PA is consistent, the Gödel sentence is true in the intended model of the natural numbers but not provable in PA.

This phenomenon is not an accident tied to a clever trick. It is structural. Whenever a theory can encode enough arithmetic to talk about proofs and computations, and its axioms are recursively enumerable, there will exist arithmetical truths that escape proof in that theory. Strengthening

PA by adding more true axioms does not solve the problem; it simply pushes the incompleteness boundary a bit higher, after which the same argument produces new unprovable truths.

For a PA-limited civilization, incompleteness is thus a fixed feature of the intellectual landscape. No amount of cleverness or incremental axiomatization will close the gap between "all arithmetical truths" and "what PA can prove." The system is not incomplete because people have not yet found the right theorems, but because the logic of self-reference and arithmetization guarantees that some truths lie beyond reach.

6.2. Examples from combinatorics and number theory

Gödel's original undecidable sentence was self-referential and deliberately constructed. Subsequent work, however, has produced natural mathematical statements—especially in combinatorics and number theory—that are true in the ordinary natural numbers but unprovable in PA.

One such example is Goodstein's theorem. Goodstein sequences are defined by a surprisingly elementary process. Start with a natural number written in base 2. Replace the base 2 by base 3 in a precise way, subtract one, then repeat with base 4, base 5, and so on. For each starting number, this procedure produces a sequence of integers that initially grows extremely rapidly. Goodstein's theorem says that, despite the explosive growth, every such sequence eventually reaches zero.

The statement "every Goodstein sequence terminates" is a simple, first-order sentence about natural numbers. It is true in the standard model of arithmetic. Yet it cannot be proved in PA. Any proof of termination requires methods that are, in a precise sense, stronger than the induction available in Peano arithmetic, typically involving transfinite induction up to certain countable ordinals.

Another example is the Paris–Harrington principle, a strengthened form of a finite Ramsey theorem. The original finite Ramsey theorem says, roughly, that for any coloring of k-element subsets of a sufficiently large finite set, one can find a large subset all of whose k-element subsets have the same color. The Paris–Harrington version adds an extra largeness condition on the homogeneous subset. This strengthened combinatorial statement is again about finite sets and natural numbers; it is true in the standard model; and it is unprovable in PA.

There are more such examples. Certain hydra game termination statements, where a simple combinatorial game always ends after finitely many steps, are true but require proof methods beyond PA. Various statements about van der Waerden-like phenomena or combinatorial well-quasi-orderings also sit above the Peano ceiling.

From the internal perspective of a PA-limited civilization, these examples would appear as conjectures with strong empirical and heuristic support, perhaps proved in stronger, unofficial meta-theories, but stubbornly resistant to any attempt at formal proof within the accepted axioms. Logicians could show that if PA is consistent, then these statements cannot be refuted in PA, and that their truth strengthens PA in specific ways. But the statements themselves would remain forever beyond official theoremhood.

The message is that incompleteness is not confined to bizarre self-referential artifacts. It touches seemingly innocent claims about the behavior of finite sequences, colorings, and games—exactly the kinds of structures that working mathematicians and computer scientists manipulate.

6.3. Program termination, safety proofs, and verification gaps

The bridge from incompleteness to program verification is short. Termination of a program on all inputs is itself an arithmetical property. For a program p, the assertion “for every input n, p halts on n” can be written as a first-order sentence over N, quantifying over encodings of inputs and computation steps.

By general arguments, one can construct programs whose termination behavior encodes deep logical facts. For example, for any consistent, reasonably strong arithmetical theory T (including PA), there exists a program p_T that halts on every input if and only if T is consistent. The construction uses arithmetization of proofs: p_T systematically searches for a proof of a contradiction in T; if it ever finds one, it enters an infinite loop; otherwise, it halts on each input after some finite but unbounded search.

For such a p_T, the statement “p_T halts on all inputs” is equivalent to “T is consistent.” If T is PA, then a PA-limited civilization cannot prove that p_PA terminates on all inputs without, in effect, proving its own consistency. Gödel’s second incompleteness theorem says that this is impossible if PA is consistent. The program p_PA may be extremely simple to describe, yet its global termination property is unprovable in PA.

This phenomenon is not confined to artificially constructed programs. Any sufficiently expressive programming language can embed logical search procedures whose halting mirrors difficult mathematical questions: the existence of counterexamples to a conjecture, the termination of other programs, or the finiteness of certain combinatorial objects. Some of these embedded questions will lie above PA’s proof-theoretic strength.

In a PA-limited world, this has direct consequences for formal verification. Many programs and systems can be proved safe and terminating using induction and standard invariants, and PA can certify those proofs. But there will also be specific programs for which no PA-proof of total correctness exists, even though meta-theoretically one might know that they are safe. Engineers

might rely on extensive testing, on proofs in stronger but unofficial logical systems, or on empirical arguments. From the official foundational point of view, however, there is a permanent gap.

These verification gaps are not merely practical limitations of time or ingenuity. They are structural: given the foundational ceiling, certain program properties cannot be certified, no matter how long one searches for proofs, unless one is willing to step outside PA. The civilization can know, in principle, that such gaps must exist, but it cannot, in general, algorithmically distinguish which specific programs fall into them.

6.4. The boundary between practical undecidability and foundational incompleteness

Not every failure to decide a mathematical or computational question is due to incompleteness. In everyday practice, people encounter problems that are undecidable in a looser sense: they are too hard, too large, or too poorly structured for current techniques. Complexity, not incompleteness, is the main obstacle. One could, in principle, settle the question within the existing axioms, but the required proof or computation is beyond available resources.

Foundational incompleteness is different. It says that even if one had unbounded time, ingenuity, and computational power, there is no proof within PA of certain true statements. The barrier is not human weakness but the logical structure of the system itself.

In real mathematical life, these two kinds of difficulty are entangled. When a problem resists solution, it is rarely clear whether it is simply hard or truly unprovable. Only in special cases, such as Goodstein's theorem or the Paris–Harrington principle, do we have meta-mathematical proofs that a statement is independent of PA and therefore sits above its ceiling.

For a PA-limited civilization, the boundary between practical undecidability and foundational incompleteness would have a complicated status. On the one hand, logicians can prove general results: that incompleteness holds, that some arithmetical truths are unprovable, that certain constructed examples sit beyond PA. On the other hand, for many naturally arising scientific and engineering questions, no one would know whether the difficulty lies in the structure of PA or merely in the current state of knowledge.

In verification, a failed proof attempt might mean that more ingenious invariants are needed, or that a different specification is required. Or it might mean that the property is in fact unprovable in PA, even though it is true. There is no general procedure that can tell the difference. The invisible ceiling is felt as a hazy horizon: known to be there in theory, but only partially mapped by specific independence results.

This ambiguity has philosophical implications. It undercuts any simple narrative in which science is imagined as marching steadily toward complete, formal mastery of its own tools. In a PA-

bounded world, even the arithmetic core of computation and discrete mathematics contains truths that remain forever beyond formal capture. Practically, the civilization may accomplish an enormous amount, often without bumping into the ceiling. Conceptually, however, it lives with the knowledge that its most basic language—arithmetic itself—harbors truths that it can recognize only from outside its own official framework.

7. Higher Set Theory and the Mathematics That Goes Missing

Peano arithmetic reaches surprisingly far into what working scientists and engineers recognize as mathematics. Yet there are entire regions of the mathematical landscape whose natural home is not arithmetic but full-blown set theory. These regions trade heavily on uncountable sets, strong compactness principles, and powerful existence axioms. They include large portions of modern measure theory, functional analysis, and higher set theory itself.

From a PA-limited perspective, much of this high structure becomes invisible or at best shadowy. One can still talk about countable approximations, definable fragments, and effective special cases, but the more sweeping theorems often rely on comprehension principles and choice axioms that go beyond what PA can justify. This section sketches what is lost or attenuated when one lives strictly under the Peano ceiling, and then asks how much that loss actually matters for empirical science.

7.1. Measure, compactness, and choice beyond PA's strength

Measure theory and topology provide some of the clearest examples of where stronger set-theoretic tools play a central role.

At the intuitive level, measure theory assigns sizes to sets of points and integrates functions over domains. For familiar, nicely behaved cases—intervals in the real line, regions with piecewise smooth boundaries—this can be recast in terms of Riemann or Lebesgue integrals over effectively describable sets. These can often be handled inside systems close to PA, using coding of sets by algorithms and restricting attention to well-presented functions.

However, full measure theory goes far beyond these tame situations. It involves σ-algebras generated by arbitrarily complex subsets of R, completion of measures with respect to sets of measure zero, and the existence of non-measurable sets that cannot be assigned a size in any consistent way if one adopts the axiom of choice. Many theorems, such as general forms of Fubini's theorem, Radon–Nikodym derivatives for arbitrary σ-finite measures, and deep structural results about measure-preserving transformations, depend on comprehension principles over sets of reals and on choice-like axioms.

Similarly, topology at its most robust involves arbitrary products, compactifications, and spaces of functions, many of which are not separable or otherwise tame. Compactness theorems that assert every open cover has a finite subcover for broad classes of spaces, or that every filter has a cluster point, often rely on principles that outstrip arithmetic. Tychonoff's theorem, for example, stating that arbitrary products of compact spaces are compact, is typically equivalent to or dependent on strong forms of the axiom of choice.

In a PA-limited world, one can still work with countable bases and second-countable spaces, such as most function spaces used in analysis and physics. Compactness in these settings often reduces to combinatorial properties of sequences and covers that can be expressed arithmetically. However, the sweeping generality of full compactness theorems and the existence of certain pathological or highly nonconstructive objects disappear from official mathematics.

The axiom of choice itself is a key dividing line. In the arithmetic setting, one can adopt countable or dependent choice schemas that are sufficient for many analytic arguments, while remaining roughly at PA strength for arithmetic statements. Full choice, which asserts that every collection of nonempty sets admits a global chooser, is much stronger and leads to consequences, such as the existence of non-measurable sets, that cannot be captured in PA. A PA-limited civilization may treat full choice as a speculative principle, useful in informal reasoning but not part of its official axiomatic base.

7.2. Functional analysis and structures requiring stronger comprehension

Functional analysis studies spaces of functions and operators acting on them. It provides the conceptual backbone for quantum mechanics, partial differential equations, and many areas of applied mathematics. At first glance, this might seem incompatible with a PA-limited foundation, but the picture is nuanced.

On the constructive and effective side, many of the Hilbert and Banach spaces used in applications are separable: they have countable dense subsets, often consisting of simple functions with rational coefficients. The operators of interest are frequently bounded, compact, or otherwise well-behaved, and can be approximated by finite-rank operators. Statements about such structures often reduce to claims about sequences of rational approximations and inequalities among them. These are amenable to arithmetical formalization.

However, the full apparatus of functional analysis reaches much further. It considers dual spaces of all continuous linear functionals, spaces of all bounded linear operators, and various completions and bidual constructions. The existence of certain linear functionals, especially those that are not given by explicit integrals or sums, is guaranteed by nonconstructive tools such as the Hahn–Banach theorem in its most general form. Many of these results are

equivalent, over weak base theories, to strong comprehension principles about sets of reals or sets of functionals.

For example, in a general Banach space, the dual may be enormous and highly nonseparable. The existence of norm-preserving extensions of functionals from a subspace to the whole space requires a form of choice or comprehension that is not justified by PA alone. The structure of C(K) spaces, where K is a compact Hausdorff space, involves rich interplay between topology and measure based on set-theoretic assumptions about K and about the underlying continuum.

In a PA-limited civilization, one can still develop a significant part of functional analysis by restricting attention to separable spaces, operators that are given by explicit formulas or effective prescriptions, and functionals that arise from integrals with respect to well-behaved measures. Much of the functional analysis used in physics, signal processing, and numerical analysis falls into this category.

What goes missing are the grand general theorems that assert existence of functionals or operators in all cases, or that classify spaces up to isomorphism based on large-scale structure. These theorems employ a level of set-theoretic comprehension beyond PA's reach. They may still be studied informally, or in stronger meta-theories, but they cannot be fully internalized as theorems of the official arithmetic foundation.

7.3. Large cardinals, forcing, and independence phenomena

Beyond the mathematics that is directly connected to analysis and functional spaces lies the realm of pure set theory: large cardinals, forcing, and independence results. These subjects live far above Peano arithmetic in strength and conceptual distance from concrete science.

Large cardinal axioms posit the existence of infinite cardinals with strong combinatorial and structural properties: inaccessible cardinals, measurable cardinals, supercompact cardinals, and many others. These axioms are not consequences of the usual Zermelo–Fraenkel axioms of set theory, but they imply rich hierarchies of sets and strong regularity properties. They are deeply tied to questions about the structure of the real line, determinacy of infinite games, and the classification of definable sets.

Forcing is a method of constructing new models of set theory by adding generic sets, often used to show that certain statements are independent of the standard axioms. This technique produced famous results such as the independence of the continuum hypothesis and of various forms of the axiom of choice. The landscape revealed by forcing and large cardinals is one in which the set-theoretic universe is not uniquely determined by basic axioms, but can be extended in many directions, each with its own combinatorial and structural features.

From the viewpoint of PA, this entire enterprise is external. Large cardinal axioms assert the existence of entities far beyond what can be coded arithmetically in any simple way. Forcing constructions manipulate transfinite hierarchies of sets and generics; they do not have straightforward arithmetical analogues. While set theorists can use PA to talk about finite approximations and combinatorial consequences of large cardinal assumptions, the core content of these theories lies outside the Peano universe.

A PA-limited civilization could still study fragments of set theory that have arithmetical consequences, such as certain combinatorial principles on the natural numbers that are implied by large cardinals. It could also reason meta-theoretically about independence, showing that some arithmetical statements cannot be settled by PA and require stronger axioms. But the full picture of multiple set-theoretic universes, each with its own continuum structure, would remain a speculative backdrop rather than part of official mathematics.

This does not mean that higher set theory is useless. It has informed our understanding of what kinds of axioms are available to extend our foundations and has clarified the limits of what can be proved in weaker systems. For a PA-limited world, however, large cardinals and forcing would be like distant mountains: visible from afar through their shadows on arithmetic, but not accessible terrain.

7.4. How little of this is needed for current empirical science

The gap between the grandeur of higher set theory and the modest requirements of empirical science is striking. Most physical theories and engineering disciplines make use of mathematical structures that can be modeled within separable spaces, countable constructions, and effective approximations. They rely on finite or countable collections of equations, boundary conditions, and constraints; on numerical methods that are explicitly implementable; and on probabilistic models that can be approximated by finite samples.

The theorems that underwrite this practice tend to live in relatively low regions of the foundational hierarchy. Existence and uniqueness theorems for differential equations under Lipschitz conditions, basic spectral results for operators on separable Hilbert spaces, convergence of numerical schemes under mesh refinement, and central limit theorems for sums of independent random variables can all be developed with modest axioms when restricted to effectively presented objects.

Very few, if any, concrete predictions in physics or engineering depend on the full generality of, say, Tychonoff's theorem for arbitrary products, the existence of non-measurable sets, or the structure of dual spaces of nonseparable Banach spaces. Even when such theorems play a role in the conceptual background or in the classification of structures, the link to experiment runs through special cases that can be described and analyzed within arithmetic frameworks.

This does not make higher set theory irrelevant. It plays a vital role in pure mathematics, in clarifying the landscape of possible universes, in showing what kinds of regularity and classification theorems are available or impossible. It also feeds back into more applied areas by suggesting new structures, techniques, and questions. But if one is concerned strictly with what is necessary to state, compute, and interpret empirical predictions at current levels of precision, the need for strong set-theoretic axioms is remarkably limited.

For a PA-limited civilization, this mismatch would be a source of both relief and humility. Relief, because it means that the foundational ceiling does not cut off access to the mathematics needed for experimental science and technology. Humility, because it underscores that even a very successful empirical science does not exhaust the higher reaches of mathematical possibility. The arithmetic foundation suffices for the practical modeling of nature, yet it leaves untouched vast regions of abstract structure that, while not needed for measurement and control, reveal something about the full space of consistent ways mathematics could extend.

In this sense, higher set theory and the mathematics that goes missing under PA are a reminder: the success of science does not depend on exploring every corner of the mathematical universe. It depends on having enough mathematics, properly aligned with the phenomena at hand. Peano arithmetic turns out to be more than enough for that purpose, even as it marks a boundary beyond which other, more speculative mathematical worlds unfold.

8. Epistemic Life Inside PA: Knowledge, Certainty, and Blind Spots

Treating Peano arithmetic as a foundational ceiling is not just a technical choice. It reshapes what a community takes itself to know, how it experiences certainty, and how it relates to truths that lie beyond formal proof. The daily life of mathematicians, scientists, and engineers in a PA-bounded culture would be structured around a sharp distinction: on one side, statements that can be encoded and proved inside the Peano framework; on the other, beliefs, conjectures, and heuristics that may be compelling but never attain the status of officially recognized theorems.

This division would not make research impossible or sterile. People already work with conjectures and informal arguments all the time. What changes is how explicitly the civilization has to live with the gap between formal certifiability and broader forms of understanding. Incompleteness makes that gap permanent. The impossibility of proving the consistency of PA blocks certain kinds of foundational reassurance. At the same time, scientific practice continues to rely on models and methods whose justification flows partly through informal channels.

This section looks at what counts as proof in such a culture, the impact of not being able to certify the consistency of one's own foundations, the role of informal knowledge, and the position of scientific realism and mathematical Platonism under PA constraints.

8.1. What counts as a proof in a PA-bounded culture

In a PA-bounded culture, the official notion of proof is sharply delimited. A proof is a finite object: a sequence of formulas in the first-order language of arithmetic, each either an instance of a Peano axiom or derived from earlier lines by an accepted rule of inference. The entire sequence is then encoded as a natural number. Whether that code represents a valid proof is itself an arithmetical property.

From the inside, this gives proofs a dual character. On the one hand, they are recognizable artifacts that mathematicians read, write, and critique in informal notation. On the other, they can in principle be fed into proof-checking software that verifies, step by step, their conformity to the Peano axioms and rules. The ideal is that every official theorem has, somewhere behind it, a proof that could be mechanically validated within PA.

This does not mean that every working mathematician writes fully formal proofs. As in our world, there would be layers of exposition: highly compressed arguments in research papers, expanded versions in textbooks, and, in principle, fully formal derivations accessible to proof assistants. What makes the culture PA-bounded is that, whenever it matters, the community is committed to the idea that these informal arguments can be unpacked and checked in Peano arithmetic.

Controversies about proofs would therefore take on a particular flavor. If there is a dispute about whether a purported theorem is genuinely established, the question reduces, at least in principle, to whether a certain arithmetical relation holds: does there exist a PA-proof of the sentence encoding the statement. In practice, no one can search all possible proofs. But the conceptual backdrop is rigid: either such a proof exists in PA or it does not. There is no higher system to appeal to as an arbiter.

At the same time, incompleteness ensures that there are true statements with no such proofs. The community can know that such statements exist in general, by meta-theoretic arguments, but it cannot always know which specific conjectures fall into that category. The official notion of proof is therefore both precise and incomplete: it is clear what it is, and also clear that it does not reach all truths expressible in its own language.

8.2. The impossibility of proving the consistency of one's own foundations

A natural desire in any foundational enterprise is to prove that the axioms one is using are free of contradiction. For a PA-bounded civilization, this desire runs directly into Gödel's second incompleteness theorem. Assuming PA is consistent, PA cannot prove a sentence that correctly expresses its own consistency. Any such proof, if it existed, would contradict the theorem.

From the inside, the community can formalize the notion of a contradiction: a proof of both a statement and its negation, or a proof of a distinguished absurdity symbol. It can write down a formula Con(PA) that says, in arithmetical terms, “there is no code of a valid PA-proof of a contradiction.” It can even reason about this formula, deriving implications such as: if Con(PA) is false, then everything is provable. What it cannot do, under its own rules, is derive Con(PA) itself.

There are two ways this might be lived.

One is a stance of pragmatic trust. The community notes that PA has withstood extensive use without yielding contradictions, that its axioms match basic intuitions about the natural numbers, and that stronger meta-theories can be used to argue for its consistency. On this view, the lack of an internal proof is an abstract limitation rather than a practical worry. People continue to work under PA much as mathematicians in our world work under set-theoretic axioms whose ultimate consistency they cannot internally prove.

The other stance is more philosophically unsettling. It emphasizes that the most basic arithmetic statements, such as “no proof of 0 = 1 exists in PA,” lie forever beyond formal endorsement by PA itself. The community must accept its own foundation on extra-formal grounds: intuition about numbers, empirical reliability of the system, or trust in meta-theories that themselves rest on unproven assumptions. The dream of a fully self-justifying foundation is decisively blocked.

In either case, the impossibility of an internal consistency proof shapes the epistemic climate. It reminds the culture that even its most elementary certainties are not grounded in a single, closed formal loop. Arithmetic is both the medium in which proofs are expressed and a subject about which there are unprovable facts. The foundation is robust enough to support science, but not transparent enough to certify itself.

8.3. Informal knowledge versus formal certifiability

Every scientific community relies on a large body of informal knowledge: heuristic arguments, computational evidence, numerical experiments, approximations, and analogies. In a PA-limited civilization, this informal layer becomes even more salient, because the line between what can be formally certified and what cannot is sharply visible and known to be incomplete.

On one side of the line are statements that have been, or could in principle be, reduced to arithmetical theorems of PA. For these, there is a clear sense in which the community can say, “we know this, in the strongest available sense.” The knowledge is anchored in a chain of reasoning that fits within a fixed deductive system whose rules are unambiguous.

On the other side lie statements that are supported by extensive evidence but resist formalization in PA. Examples include conjectures in number theory that have been verified for trillions of cases, combinatorial principles known to have strong consequences, or termination assertions about programs whose behavior has been extensively tested and is backed by higher-level arguments outside PA. For these, the community may have a high degree of confidence, perhaps approaching certainty in practice, yet without PA-provable status.

The result is a multi-layered epistemic ecology. Scientists and engineers routinely rely on informal knowledge when designing systems or interpreting data, while reserving formal certifiability for parts of their work where absolute guarantees are required. Pure mathematicians may distinguish between theorems officially in PA, theorems provable in stronger but unofficial frameworks, and conjectures supported by heuristic and computational arguments.

Incompleteness gives this ecology a structural justification. It is not merely that formalization is hard or expensive; it is that some truths cannot be brought within PA at all. The presence of such truths legitimizes a category of well-motivated beliefs that permanently lack the status of PA theorems. A PA-bounded culture therefore lives with a principled pluralism of justification: formal proof where possible and needed, plus a spectrum of extra-formal reasons elsewhere.

8.4. Scientific realism and mathematical Platonism under PA constraints

Finally, how does a PA-limited foundation bear on philosophical stances about science and mathematics themselves?

For scientific realism, the key question is whether the world described by physical theories is taken to exist independently of our descriptions. In a PA-bounded culture, there is no intrinsic conflict between realism and arithmetic foundations. The physical world can be thought of as whatever gives rise to regularities that can be modeled by discretized equations and approximated solutions. The fact that the official mathematics sits inside PA does not prevent scientists from believing that they are discovering objective structures in nature; it simply means that their verified models are expressed in a particular formal language.

Indeed, the emphasis on effective approximations and finite simulations may push scientific realism toward a modest, operational form. Rather than hypothesizing grand structures that require strong set-theoretic axioms to describe, the culture centers its realism on what can be captured by algorithms and countable models. The success of these models in prediction and control becomes the main warrant for believing in an external world with certain patterns.

For mathematical Platonism, the situation is more delicate. Platonism, in one form, holds that mathematical objects such as numbers, sets, and functions exist independently of human thought and that mathematical statements are true or false about this realm. A PA-bounded

foundation does not rule out such a view, but it forces a distinction between two levels of mathematical reality.

One level is the arithmetic universe itself: the natural numbers and their properties. A PA-limited Platonist might hold that there is a single, determinate structure of N about which statements are objectively true or false, including those that PA cannot prove. On this view, PA is a partial window into a richer arithmetical reality. True but unprovable statements are real facts about N that our chosen system simply cannot access from within.

A second, more expansive level would include higher set-theoretic universes: the full power set of the reals, large cardinals, and so on. For a PA-bounded civilization, this realm is accessible only via speculative meta-theories. Some mathematicians may adopt a Platonist stance toward these higher entities, believing that they exist and that statements about them have definite truth values, even though such beliefs play little role in empirical science. Others may be more cautious, treating higher set theory as a conceptual playground whose structures do not demand the same kind of ontological commitment.

The constraint to PA as an official foundation therefore does not settle philosophical debates. It does, however, sharpen them. Realists and Platonists must articulate how they relate to truths that are known, by meta-theoretic argument, to be beyond PA. Formalists and structuralists, in turn, can point to the success of PA-based mathematics in science as evidence that much of what matters can be captured without invoking a maximal mathematical universe.

In all these discussions, the central fact remains: Peano arithmetic is both expansive and incomplete. A PA-bounded civilization can build rockets, design computers, and develop sophisticated theories of nature, all while acknowledging that its most basic mathematical framework cannot see every truth even about the natural numbers. Its epistemic life unfolds under a ceiling that is invisible in daily practice yet provably there, shaping the contours of what can be formally known and what must remain, at least from the inside, forever conjectural.

9. Comparing PA-World Science with Our Own

A civilization bounded by Peano arithmetic is not a civilization of diminished intelligence or curiosity. Its scientists still observe, model, and manipulate the world; its engineers still build devices; its mathematicians still push at the edges of what can be proved. The difference lies in what is officially countenanced as a proof and in how far the formal layer of their intellectual life extends.

Viewed from outside, such a civilization would look strangely familiar. Laboratories, simulations, and devices would not obviously betray the foundational constraint. The divergences would

show up more in the upper atmosphere of pure mathematics and in the background assumptions of education and philosophy than in the day-to-day conduct of empirical work.

9.1. Would experimental and technological trajectories differ?

For most of the technological history we can presently imagine, the short answer is: probably not much.

Experimental science is rooted in observation, measurement, and model-building. Models are compared with data; their parameters are adjusted; their predictions are tested. This cycle depends heavily on statistics, numerical analysis, and differential equations, all of which have effective, finitary cores that fit comfortably inside PA. When one simulates a fluid, designs a control system, or trains a machine learning model, the actual computations are finite and discrete.

In a PA-world, the mathematical underpinnings of these practices would be framed more explicitly in arithmetic. Error bounds for numerical schemes, stability proofs for control systems, and convergence guarantees for algorithms would be expressed as statements about natural numbers representing grid sizes, time steps, sample sizes, and iteration counts. But the physics would still be modeled, the codes would still run, and the devices would still work.

Where trajectories might diverge is in areas where highly nonconstructive existence results are used as a source of confidence or inspiration. For example, certain sophisticated results in ergodic theory or functional analysis rely on strong forms of compactness or choice. In our world, these theorems sometimes provide conceptual justification for treating certain statistical behaviors as generic or typical. A PA-limited civilization might be more cautious about invoking such generalities, preferring statements that can be recast as effective approximations.

Even here, however, the impact on concrete technology is likely modest. Engineers and applied scientists already favor explicit constructions and computable estimates. The thrust of technological development is driven more by what can be simulated, controlled, and manufactured than by the existence of nonconstructive objects. A PA-bound on formal foundations would reinforce this tendency rather than curtail it.

The most palpable difference would be attitudinal. In safety-critical domains, formal verification of software and hardware properties would be explicitly understood to live inside PA. Engineers would know that some properties cannot be certified there, and this awareness might influence design choices: systems might be structured to avoid relying on behavior that appears, even in principle, to sit near the incompleteness boundary.

9.2. Where frontier pure mathematics diverges

The more dramatic divergence lies in pure mathematics at the frontier.

In our world, a substantial part of high-level research in set theory, descriptive set theory, and parts of functional analysis is explicitly set-theoretic. It invokes large cardinals, forcing extensions, and independence results. Even outside set theory proper, some classification theorems and structural results quietly depend on choice or other strong axioms.

In a PA-bounded culture, that entire region becomes, at best, unofficial. Mathematicians might still study set-theoretic universes as abstract objects in a meta-theory, but the consequences of that work would not automatically feed into the official body of theorems. Priority would go to subjects whose results can be traced back to arithmetic proofs.

This does not mean that pure mathematics collapses into elementary number theory. Instead, research priorities tilt toward areas where the boundary of PA is interesting and tractable:

- finitary combinatorics and its links to proof theory
- computable analysis and effective versions of classical theorems
- classification problems where structures are presented by algorithms or explicit generators
- reverse mathematics, mapping exactly which principles are needed for which theorems

The study of strong principles does not disappear; it changes character. Rather than being embraced as new axioms, they are treated as hypotheses whose arithmetical consequences are examined. A mathematician might say, for example: “Assuming a certain combinatorial schema not provable in PA, we can derive these finite consequences,” and then carefully analyze which of those consequences already lie below the Peano ceiling.

Some areas that, in our world, lean heavily on general set-theoretic machinery would likely be slower to develop or would develop in a more restricted, effective form. Certain classification theorems for operator algebras, measure algebras, or complex Banach spaces might be considered aspirational rather than canonical, their full generality never entering the official canon.

9.3. Possible long-term effects on education and conceptual imagination

Foundations filter into education. Over time, a PA-bounded framework would shape how students are taught to think about mathematics and its relation to the world.

Early education would still introduce real numbers, limits, and continuous functions, but with a more explicit emphasis on approximations and algorithms. The idea that a real number is

something one can approximate to arbitrary precision by a computable process might be presented not as an optional viewpoint but as the default. Infinite sets and uncountable spaces would be mentioned, if at all, as idealizations that must be grounded in discrete encodings.

This emphasis could encourage an algorithmic imagination. Students might be trained to ask, almost reflexively: "How would I compute this? How is this object given to me?" Proofs that rely on informal appeals to "all subsets" or "arbitrary functions" might be treated with caution, unless they can be unpacked into arithmetical arguments after the fact.

The downside is that some kinds of conceptual breadth might be less cultivated. The notion of a vast set-theoretic universe populated by hierarchies of infinities would be less central. Intuitions tied to uncountable combinatorics, or to large-cardinal-inspired regularity phenomena, might be rare. Students might not be encouraged to imagine spaces of all possible structures of a certain type when those structures cannot be effectively presented.

The upside is a tighter coupling between intuition and computability. The mental habits of the culture would lean toward constructive and semi-constructive reasoning, toward seeing mathematical objects as processes or codes. This might make transitions into areas like computer science, proof verification, and numerical modeling smoother, at the potential expense of a looser, more "set-theoretic" style of imagination.

In physics and engineering education, the PA-bound would mostly reinforce tendencies already present: heavy use of discretized models, emphasis on numerical techniques, and a pragmatic approach to existence theorems. The difference is that the foundational picture supporting this pedagogy would be explicit and unified, rather than implicit and heterogeneous.

9.4. Hybrid strategies: unofficial stronger reasoning, official PA bookkeeping

Of course, human practice is rarely as clean as foundational thought experiments suggest. Even in a PA-limited civilization, there would be strong incentives to develop and use richer logical tools at the informal level.

One plausible pattern is a hybrid strategy:

- At the exploratory and heuristic level, mathematicians and scientists allow themselves to reason in stronger systems: second-order arithmetic, set theory, or even large-cardinal frameworks. These tools help generate conjectures, suggest structures, and guide intuitions.
- At the certification level, when a result matters for safety, technology, or the official corpus of mathematics, its statement is translated into arithmetic and its proof is checked against PA.

Something like this already happens in our world. Many results are first discovered with the help of informal reasoning and only later, if at all, made precise in a formal system or checked by automated provers. The PA-bound thought experiment makes that division of labor explicit and principled. Strong systems are treated as external notebooks; PA is the ledger of record.

This hybrid practice has several advantages. It preserves the creative freedom of working mathematicians, who can still explore speculative axioms and abstract universes, while maintaining a clear standard for what counts as firmly established. It also encourages the development of translation techniques: methods for taking proofs in stronger systems and extracting their PA-provable arithmetical consequences.

Over time, this might lead to a refined art of “proof mining” and “foundations-aware” modeling. Researchers would become adept at asking, of any argument, which parts genuinely require strong principles and which can be pulled down into arithmetic. The invisible ceiling of PA would not be an absolute barrier so much as a contour around which mathematical practice flows, with stronger reasoning above it and arithmetical bookkeeping below.

In that sense, the contrast between PA-world science and our own is not absolute. It is more a matter of emphasis and self-awareness. Our world already lives with multiple foundations and with a gap between informal reasoning and formal verification. A PA-bounded civilization takes one of those foundations, Peano arithmetic, and elevates it to the unique official standard. The scientific and technological trajectories remain broadly similar; what changes is the clarity with which the limits of formal certainty are acknowledged and the ingenuity with which practice adapts to them.

10. Broader Philosophical Implications

The thought experiment of a civilization bounded by Peano arithmetic begins as a technical exercise in logic and foundations. But its consequences spill into broader philosophical questions: what it means to formalize, what arithmetic captures and leaves out, how a scientific culture should think about completeness, and how we ought to view current efforts to mechanize reasoning.

The picture that emerges is neither triumphalist nor despairing. Peano arithmetic is powerful enough to support an advanced technological civilization; it is also provably unable to exhaust all arithmetical truths. The lesson is not that formalization has failed, but that its limits are themselves part of what we have learned about rationality. The ceiling is not a flaw in the building; it is one of its most carefully charted features.

10.1. Limits of formalization as a feature, not a bug

It is tempting to treat the incompleteness of arithmetic as a disappointment, a sign that the dream of fully capturing mathematics in a neat axiomatic package has broken down. From this angle, Gödel's theorems and the unprovable truths they reveal look like bugs in the system: regrettable constraints on what logic can do.

The PA-world thought experiment suggests a different attitude. A civilization that works under Peano arithmetic does not grind to a halt. It builds theories, proves theorems, constructs devices, and explores nature. It also discovers, as a theorem of its own meta-theory, that there are truths about its basic subject matter that it cannot prove. This discovery is itself a triumph of formal reasoning. It exposes a structural property of any effective axiomatization of arithmetic, revealing a genuine feature of the landscape rather than a failure of the tools.

The limits of formalization, in this sense, are part of what formalization teaches. Trying to codify proof, define computability, and axiomatize arithmetic leads naturally to results about what cannot be done within any such codification. The incompleteness phenomenon is not an external constraint imposed on logic from outside; it is a pattern that emerges when logic turns its tools on itself.

Treating these limits as features has two important consequences. First, it reframes expectations. Instead of seeking a perfect system that captures all truths, we recognize that any workable system will leave some truths out and that this is a mathematically demonstrable fact, not a remediable defect. Second, it suggests a more dynamic picture of foundations. Rather than a final resting place, a system like PA becomes a well-understood base camp: secure enough for most purposes, but known to sit below higher peaks that one can only glimpse indirectly.

10.2. Arithmetic as a window onto but not a capture of mathematical truth

Peano arithmetic has a special status. It is not just one formal system among many; it is the canonical attempt to axiomatize the basic properties of the natural numbers. The PA-world scenario makes clear how much of mathematics and science can be routed through this single window. At the same time, incompleteness shows that the window, though broad, does not span the entire horizon.

From one perspective, arithmetic looks like a universal medium. Almost any discrete structure can be coded by natural numbers; almost any finite or countable mathematical construction can be represented in terms of arithmetical relations among codes. Computation, discrete probability, combinatorics, and much of analysis and physics can be framed in this language. If one were to choose a single domain as the backbone of formal reasoning for a scientific civilization, N would be a natural choice.

Yet the existence of true but unprovable arithmetical statements shows that arithmetic as formalized in PA is not the whole story. There is a distinction between the intended structure of the natural numbers and the theory we use to reason about it. The former has a full truth value for every sentence in the language of arithmetic; the latter captures only a subset of those truths. Arithmetic is thus both a subject and a lens: it is what we are talking about, and the medium through which the rest of our formal discourse is filtered.

Philosophically, this invites a layered view of mathematical truth. At the innermost layer is the structure of N itself, with its truths and falsehoods independent of any axioms we adopt. At the next layer are formal systems like PA that pick out some of those truths as theorems. At higher layers are stronger systems that pick out more truths, especially those with complex combinatorial or set-theoretic content. No single layer, if it is effective and consistent, can encompass all truths at once.

For a PA-bounded civilization, arithmetic is thus a window onto a wider reality, not a capture of it. The civilization sees enough to do science, build technology, and develop a rich corpus of mathematics, but it also comes to know that there are truths beyond the reach of its chosen pane of glass. The limitation is not on the reality of N, but on what any one syntactic apparatus can extract from it.

10.3. The meaning of “completeness” for a scientific civilization

Completeness is a slippery concept. One might hope that a scientific civilization could, in principle, arrive at a complete theory of its subject matter, one that settles every question expressible in its language. In physics, this hope usually targets the laws of nature; in logic, it targets the underlying deductive system. The PA-world thought experiment exposes tensions in this ideal.

On the logical side, completeness in the strongest sense is ruled out. For any effective, consistent system that encompasses Peano arithmetic, there will be arithmetical statements it cannot decide. A civilization that wants its proofs to be checkable by mechanical means must accept that some questions, even purely mathematical ones, will remain forever open within its official framework.

This does not mean that science cannot be complete in more modest senses. A theory might be empirically complete for a given domain if it manages to predict and explain all observable regularities at the relevant scales. A mathematical theory might be effectively complete for a given class of problems if, whenever a statement in that class is true, there is an algorithmically obtainable proof within the theory. These local notions of completeness do not conflict with incompleteness theorems, because they apply to restricted families of statements and to practical contexts.

The PA-world suggests that a scientific civilization should treat completeness as domain-relative and scale-sensitive. It can reasonably aspire to complete coverage of certain classes of phenomena or problems, while acknowledging that the total space of expressible questions is too large and structurally too intricate to be exhausted. In this picture, incompleteness becomes a boundary condition on rational inquiry: a reminder that there will always be more truths available than any one formal apparatus can domesticate.

The meaning of completeness then shifts from "having all the truths" to "having enough." Enough to make reliable predictions; enough to support robust technology; enough to organize and compress experience. A PA-based foundation meets this criterion for much of what science cares about, even while leaving some corners of arithmetic in permanent twilight.

10.4. Lessons for current debates on foundations and mechanized reasoning

Finally, the PA-world thought experiment offers perspective on contemporary debates about foundations and mechanized reasoning, including formal proof assistants and AI systems that manipulate symbols.

One lesson is about humility. The existence of a strong, well-understood system like PA that is both enormously powerful and provably incomplete should make us cautious about claims that any particular formalism, however sophisticated, will fully capture mathematical or scientific reasoning. Mechanized tools can check proofs, explore spaces of possibilities, and sometimes discover new arguments, but they operate within some fixed axiomatic and logical framework. If that framework interprets arithmetic, it will inherit the same incompleteness constraints.

Another lesson is about where the real constraints lie. The PA-world shows that for many practical purposes—engineering verification, algorithm analysis, numerical modeling—our current or slightly strengthened formalisms are more than adequate. The bottlenecks are often human creativity, computational resources, and modeling choices, not foundational strength. That suggests directing energy toward tools that help with formulation, search, and explanation, rather than toward ever-stronger axioms aimed at elusive global completeness.

At the same time, the thought experiment underlines the value of foundational diversity. Different systems illuminate different aspects of mathematics. Peano arithmetic highlights finitistic structure and effective content; second-order arithmetic and set theory highlight more global regularities and classification results. Mechanized reasoning can benefit from this diversity by allowing multiple logics and axiom sets, tuned to different tasks, rather than insisting on a single monolithic foundation.

Finally, the PA-world perspective fits naturally with a centaur-like vision of reasoning, in which humans and machines collaborate. Machines excel at checking, exploring, and managing complex formal objects; humans supply conceptual leaps, informal insight, and judgments

about which questions matter. Incompleteness guarantees that there will always be a space of insight that cannot be fully subsumed by any fixed formal engine. The role of mechanized reasoning is not to replace that insight, but to extend and discipline it within known limits.

In that sense, a civilization that knowingly lives inside Peano arithmetic may be better prepared for the age of mechanized reasoning than one that secretly hopes for a final, all-encompassing formalism. It understands that any such system, however powerful, will both enable and constrain; that proof assistants can certify much but not everything; and that the most enduring achievement of formalization is not capturing all truth, but revealing, with precision, where and why truth escapes capture.

11. Conclusion

The thought experiment of a civilization bounded by Peano arithmetic began with a simple inversion. Instead of treating PA as a low rung on a ladder toward richer foundations, we imagined it as the ceiling: the single official framework within which all formal proof must fit. Following that assumption across computation, mathematics, physics, engineering, and philosophy reveals a landscape that is neither impoverished nor complete. Peano arithmetic is far more capable than one might first suspect, yet it is also provably unable to capture all truths even about its own basic domain.

The result is a picture of scientific rationality that is both robust and humbled. A PA-bounded civilization can pursue advanced technology and sophisticated theory, but it must do so with a clear understanding that some true statements, including some about the very algorithms and structures it uses, cannot be certified within its official language. The ceiling is real, but much higher than everyday practice ever touches.

11.1. How much can be accomplished inside PA

From a practical standpoint, the answer is: almost everything that presently passes for ordinary science and engineering.

Peano arithmetic can encode Turing machines, programs, and finite data structures. It can express and prove correctness for a wide class of algorithms, describe complexity bounds, and formalize the core theorems of computability. It can reconstruct integers, rationals, and reals via arithmetical codings, and develop large parts of real analysis and calculus, at least for effectively described objects and spaces. It can express and prove a substantial body of results in finite combinatorics, graph theory, and algebra over countable structures.

In physics and engineering, the situation is similar. Models based on differential equations can be given arithmetic clothes via discrete approximations and effective existence theorems.

Stability and control can be treated as inequalities over finite sequences and matrices. Quantum mechanics on separable spaces can be recast in terms of finite approximations to operators and states. Numerical simulation, the workhorse of modern applied science, is almost tailor-made for PA: it consists of finite computations over coded arrays of rationals, accompanied by arithmetical error bounds.

In short, a PA-bounded civilization can design computers, launch spacecraft, develop quantum devices, run climate models, and build communication networks, all with formally checkable mathematics. The sheer scale of what is achievable within this framework undercuts any easy assumption that one needs vast set-theoretic universes to support advanced science. For the concrete tasks of modeling, predicting, and controlling physical systems at accessible scales, arithmetic foundations suffice.

11.2. Where the arithmetic ceiling actually matters

The real constraints show up not in day-to-day practice but at the margins: in frontier pure mathematics, in global properties of algorithms, and in the philosophical self-understanding of the civilization.

In mathematics, the ceiling matters when one moves beyond effectively presented structures to full generality. Strong compactness principles, unrestricted measure theory, nonseparable functional analysis, and higher set theory all rely on comprehension axioms and choice principles that exceed PA's strength. The existence of certain functionals, operators, and sets cannot be guaranteed within Peano arithmetic alone. Theorems about arbitrary products, arbitrary covers, or arbitrary measurable sets fall into a region where PA must either restrict their scope or regard them as external hypotheses.

In combinatorics and number theory, the ceiling appears in the form of true but unprovable statements. Goodstein's theorem, strengthened finite Ramsey principles, and hydra-like termination results are examples of simple-looking arithmetical sentences that cannot be proved in PA despite being true in the standard natural numbers. These are not curiosities but markers that even purely finitary reasoning can reach beyond the system's proof-theoretic reach.

In computation and verification, the ceiling manifests in the impossibility of certifying certain termination and safety properties. For some programs, the statement that they halt on all inputs is equivalent to or stronger than the consistency of PA itself. No PA-proof of such termination can exist if the system is consistent. This creates structural verification gaps: regions where the behavior of a program may in fact be entirely benign, yet cannot receive formal endorsement within the official framework.

Finally, at the foundational level, the ceiling blocks any internal proof of PA's own consistency. A PA-bounded civilization must rely on extra-formal judgment, or on stronger meta-theories, to justify its trust in its basic axioms. It can prove, within PA, that if a contradiction were provable then everything would follow, but it cannot, on pain of inconsistency, prove that no such contradiction exists.

These are not practical showstoppers. Most research never approaches these boundaries. But they matter for how the civilization understands the scope and limits of its own rational tools.

11.3. Open questions for logic, philosophy, and the practice of science

The PA-world thought experiment raises a family of open questions that cut across disciplines.

For logic, one question is how far one can push the extraction of PA-provable content from proofs carried out in stronger systems. Given a theorem proved using, say, second-order arithmetic or set theory, to what extent can its essential arithmetical core be captured in PA? This is related to proof mining and reverse mathematics: the systematic study of what axioms are really needed to prove which theorems, and how those proofs can be transformed.

Another logical question concerns the structure of unprovable truths. Beyond a few celebrated examples, how dense are the independence phenomena in everyday mathematics? How often do natural conjectures encountered in computation, combinatorics, or dynamics turn out to sit above PA? Mapping this boundary more clearly would give a more detailed picture of the arithmetic ceiling's shape.

For philosophy, the thought experiment invites reflection on how a scientific culture should think about truth, proof, and rational belief in the presence of known incompleteness. What epistemic status should be accorded to statements that are strongly supported by evidence and higher-level reasoning but known to be unprovable in the official system? How should one relate to axioms that one cannot, by design, fully justify from within?

For the practice of science, an open question is how explicit awareness of foundational limits should influence methodological choices. Should engineers systematically design systems to avoid specifications that appear to border on incompleteness phenomena? Should physicists favor models whose key predictions can be backed by proofs in weak systems, or is that an unnecessary constraint? As formal verification and mechanized reasoning become more widespread, these questions become less abstract and more practical.

11.4. Final reflections on living with necessary incompleteness

The deepest lesson of the PA-world is that incompleteness is not an unfortunate accident at the fringes of reasoning, but a structural aspect of any attempt to formalize reasoning about basic arithmetic. A scientific civilization that anchors its formal life in Peano arithmetic can accomplish

a tremendous amount. It can probe the cosmos, engineer complex technologies, and develop rich theories. The limits it encounters are subtle: they do not block rockets from flying or codes from running, but they do block certain kinds of ultimate assurance.

Living with necessary incompleteness means acknowledging that there is no final, self-contained formal theory that captures all truths even in as narrow a domain as the natural numbers. It means accepting that proof, as mechanizable derivation from fixed axioms, is both central and bounded. It also means recognizing that informal reasoning, empirical evidence, and higher-level mathematical frameworks will always play a role alongside any fixed formal system.

For our own world, which does not in fact restrict itself to PA but lives amidst multiple overlapping foundations, the PA-bound thought experiment offers perspective. It shows that one can be both rigorous and limited, both technologically advanced and formally incomplete. It suggests that the real choice is not between full capture and failure, but between being explicit about our ceilings and pretending they are not there.

A civilization that grows up under Peano arithmetic would, if it is reflective, learn to see its foundation as a powerful instrument rather than a final authority. It would treat formalization as a way of organizing and checking certain kinds of reasoning, not as the sole arbiter of truth. It would cultivate tools and practices that respect the arithmetic ceiling while exploring, conceptually and meta-theoretically, what lies beyond it.

In that sense, the lesson is not that we should shrink our ambitions to fit PA, nor that we should abandon formalization in the face of incompleteness. It is that mature rationality involves working energetically within well-understood limits, while keeping an honest eye on the truths that no single formal system can ever fully bring inside.

12. References

Bernays, P., & Hilbert, D. (1934–1939). *Grundlagen der Mathematik* (Vols. I–II). Springer. Internet Archive+1

Buss, S. R. (Ed.). (1998). *Handbook of Proof Theory*. Studies in Logic and the Foundations of Mathematics, Vol. 137. Elsevier. MathWeb+1

Church, A. (1936). An unsolvable problem of elementary number theory. *American Journal of Mathematics*, 58(2), 345–363. UC Irvine ICS+1

Feferman, S. (1988). Hilbert's program relativized: Proof-theoretical and foundational reductions. *Journal of Symbolic Logic*, 53(2), 364–384. Cambridge University Press & Assessment+1

Gentzen, G. (1936). Die Widerspruchsfreiheit der reinen Zahlentheorie. *Mathematische Annalen*, 112, 493–565. Wikipedia

Gödel, K. (1931). Über formal unentscheidbare Sätze der *Principia Mathematica* und verwandter Systeme I. *Monatshefte für Mathematik und Physik*, 38, 173–198. SCIRP+1

Goodstein, R. L. (1944). On the restricted ordinal theorem. *Journal of Symbolic Logic*, 9(2), 33–41. FA EWI TU Delft+1

Nelson, E. (1986). *Predicative Arithmetic*. Princeton University Press. Princeton Math+1

Paris, J., & Harrington, L. (1977). A mathematical incompleteness in Peano arithmetic. In J. Barwise (Ed.), *Handbook of Mathematical Logic* (pp. 1133–1142). North-Holland. PhilPapers+1

Peano, G. (1889). *Arithmetices principia, nova methodo exposita*. Torino: Fratres Bocca. GitHub+3Wikipedia+3Internet Archive+3

Simpson, S. G. (2009). *Subsystems of Second Order Arithmetic* (2nd ed.). Cambridge University Press. sgslogic.net+2aslonline.org+2

Turing, A. M. (1936). On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42(1), 230–265. Virginia Computer Science+2londmathsoc.onlinelibrary.wiley.com+2

Weiermann, A. (2005). Analytic combinatorics, proof-theoretic ordinals, and phase transitions for independence results. *Annals of Pure and Applied Logic*, 136(1), 189–218. PhilPapers+1

(Additional survey and expository sources on Gödel's incompleteness theorems, Gentzen-style proof theory, reverse mathematics, and Paris–Harrington–type independence results may be appended here in a final version, depending on editorial preferences.)

The author has published over 4,600 AI-assisted papers at: https://www.researchgate.net/profile/Douglas-Youvan/research . Google Scholar has indexed these preprints, and if you want a particular reference, the AI mode of a Google search (Gemini) has efficiently catalogued these preprints.